IETF Behave Working Group C. Perkins
Internet-Draft WiChorus Inc.
Expires: September 9, 2010 X. Li
C. Bao
CERNET Center/Tsinghua University
March 8, 2010

Comparison and integration of IVI, IVI+DHCP, 1:N IVI and SIPNAT
draft-nat46compare-perkins-01.txt

Abstract

IVI, IVI + DHCPv6, 1:N IVI, Bidirectional NAT, and SIPNAT have been proposed to global Internet access to IPv6-only domains. These methods are not all mutually exclusive, and we propose that IVI (along with its variants) along with SIPNAT may be considered together as a more complete solution for enabling access from today's IPv4 global Internet into IPv6-only domains. This would likely have the effect of accelerating the adoption of IPv6, since with such access enablements there would be much reduced incentive for new deployments to require IPv4 addressability. As part of the discussion, we also include a comparison of the various techniques so the relative trade-offs may be more easily understood.

1. Introduction

IVI [2], 1:N IVI [4], and SIPNAT [1] have been proposed to enable communications to be initiated from today's IPv4-based Internet and successfully terminated at IPv6-only devices. This is important so that IPv6-only devices will be fully functional and reachable in today's IPv4 Internet.

In particular, it is not enough to only support outgoing connections (i.e., connections initiated by the IPv6 network device not directly reachable from today's IPv4 Internet). For many purposes, it is required to support incoming connections. Businesses need to serve communications from their clients, which almost always seem to emanate from the global IPv4 Internet. This means that such communications have to be admissible even when the target computer (in the internal domain) has not triggered the setup operations at the border router which are typically required for operations such as NAT to work. Such communications, which are initiated from sites external to the site hosting the IPv6-only destination are classified as belonging to "Scenario 2" or "Scenario 4" [3]

With traditional port-mapped NAT (NAPT), this has not been possible because, for each source-destination flow, the translation parameters for the flow have had to be established by the internal network node (i.e., the node with the IP address that is incompatible with the addressing domain of the global Internet). In particular, for each such flow there needs to be an external IP address and an external port assigned. Packets arriving at the external IP address and port are then translated and retransmitted with new IP headers containing the translated IP address and port number. This works for IPv6-->IPv4 translation, IPv4-->IPv4 translation (e.g., today's Internet), and other variations as well. It is a workable solution (with various second-order difficulties) for enabling outgoing traffic to be delivered into the global Internet.

But any business requires global presence and continuous, on-demand availability. The customers have to be able to initiate contact with the business services, not the other way around. Similarly for all other online service organizations (including governmental, non-profit, and family websites).

2. Proposals for supporting Scenario 2 and Scenario 4 communications

To enable such incoming translations IVI [2] (along with several variations for improved scalability) and "source-IP NAT" (SIPNAT)[1] have been proposed. It is not the purpose of this document to reiterate the individual designs. Instead, we exhibit the characteristics of each proposal in such a way that it is easier to identify the best translation technique according to the needs of the specific IPv6-only destinations. For some destinations that experience persistent high-volume traffic, IVI is best. For some destinations that support significant traffic which is variable from time to time, IVI+DHCP is best. For destinations that host certain appropriate applications, 1:N IVI may be best. For some other destinations hosting relatively lower-volume websites, SIPNAT may be best.

IVI

IVI [2] statically reserves an IPv4 address on the translator for each IPv6 destination which is to be made reachable to the global IPv4 Internet. This is done via a one-to-one mapping algorithm between IPv4 and IPv6 addresses. The IPv4 address can only be used by the single IPv6 destination.

IVI + DHCP

(IVI + DHCP) uses DHCP to dynamically reserve an IPv4 address on the translator, and as part of the reservation process updates the DNS so that the allocated address is returned for the IPv6 destination to be made reachable. At any one time, the IPv4 address can only be used by the single IPv6 destination, but the IPv4 address can be reused for another IPv6 destination depending on the DHCP lifetime allocation policies.

1:N IVI

1:N IVI [4] statically reserves an IPv4 address and a port range (either high or low) on the translator for each IPv6 destination which is to be made reachable to the global IPv4 Internet. For example, this could be done by preconfiguration The IPv4 address and port range and can only be used by the single IPv6 destination. DHCP could be extended for this purpose to also return the port range when the IPv4 reservation is made. This would allow temporal sharing for the address and port range.

SIPNAT

SIPNAT ("Source-IP NAT") dynamically associates an IPv4 address on the translator, on demand, when a communication is initiated from the global IPv4 Internet. The FQDN of the destination IPv6 does not typically have a persistent resolution for any particular IPv4 address. A single IPv4 address of the translator can be simultaneously used for flows to many different IPv6 destinations. The source IPv4 address of incoming packets is used to identify the desired IPv6 destination. For any such source IPv4 address, only one destination is reachable at the translator's IPv4 interface address which as been associated with the IPv6 destination. Ports are not required for the translation, but should be considered as part of of the set of flow translation parameters.

The authors believe that all of these approaches have interesting use cases and should be coordinated into an overall solution.

- IVI is a good solution for destinations that serve many flows, and are each typically active serving one or more flows from the IPv4 Internet. In this case, there is little or no advantage to be gained by dynamic assignment.

- IVI + DHCP is a good solution for destinations that may be inactive for extended periods of time, but are also likely to serve many flows during other extended periods of time.

- 1:N IVI is a good solution which may be workable for situations where destinations host applications that are resilient and aware of port restrictions.

- SIPNAT is a good solution for situations where there are a large number of IPv6 destinations, each of which serves a relatively low volume of flows (e.g, fewer than 100 distinct flows per hour).

3. Work to be done

The usefulness of this integrated solution set depends upon establishing policies for partitioning the available IPv4 addresses of the translator according to persistence of IPv4 address and the persistence of the FQDN resolution. For sites requiring full persistence, IVI should be used. Sites serving a continual stream of new flow requests fit within this classification. For sites requiring persistence of association between FQDN and IPv4 address, (IVI + DHCP) should be used. Sites which sometimes serve a high volume of flow requests, but at other times are quiescent, fit this second classification. Sites which serve flows less frequently (e.g., approximately 1 per minute or so) can be served by SIPNAT.

The exact percentages and flow rates for these classifications remain to be determined. As experienced is acquired in the management of larger domains, better estimates for economical allocations of IPv4 addresses will become feasible. In general, the more IPv4 addresses that are available on the translator, the better. But, on the other hand, the fewer IPv4 addresses required, the better. Soon, the IPv4 address space will be exhausted and it is important to improve as much as possible the utilization of this valuable resource.

The partitioning of the translator's IPv4 address space is to be made between IVI, (IVI+DHCP), and SIPNAT. This partitioning may be static, based on observed traffic characteristics. As our understanding improves, we suggest that the partitioning itself should be made dynamic. For instance, if a particular website grows in popularity so that it is constantly serving new flow requests, it may be advisable to remove it from SIPNAT handling by assigning a permanent IVI address on the NAT IPv4 interface. Conversely, a particular IVI website may be partitioned into multiple destinations, each of which might have lower traffic volume and therefore require a lower persistence of addressability, potentially served well by (IVI + DHCP) or SIPNAT. Many variations are likely to be found useful.

4. Comparison Chart

The comparison chart is made based on the following characteristics:

Translator complexity

Stateless

DNS decoupled

Conserves IPv4 addresses

Application Port Transparency

Continuous service (vs. dynamic assignment)

	Stateless	Decoupled	Conserves	App Transp.	Port xlate	Cont.
1:1 IVI	Y	Y	N	Y	N	Y
1:1 IVI+DHCP	Y	Y	Y	Y	N	N
1:N IVI	Y	Y	Y	N	Y	Y
Bidirectional NAPT	N	N	Y	N	Y	N
SIPNAT	N	N	Y	Y	N	N

where:
"Decoupled" means "Decoupled from DNS"
"Conserves" means "Conserves IPv4 Addresses"
"App Transp." means "Application Transparency"
"Port xlate" means "Ports are translated"
"Cont." means "Continuous service at all times"

Figure 1: Comparison of approaches

5. Summary

Several solutions have been proposed for the cases denoted as "Scenario 2" and "Scenario 4". We have found that these solutions can be coordinated as described in this document, and propose that our solutions be considered together for standardization in fulfillment of the requirement.

6. References

6.1. Normative References

[1] Perkins, C., "Translating IPv4 to IPv6 based on source IPv4 address", draft-perkins-sourceipnat-01 (work in progress), October 2009.

[2] Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The CERNET IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition", draft-xli-behave-ivi-02 (work in progress), June 2009.

[3] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", draft-ietf-behave-v6v4-framework-07 (work in progress), February 2010.

[4] Li, X. and C. Bao, "Stateless/Partial-state 1:N Network Address and Protocol Translation between IPv4 and IPv6 nodes", draft-xli-behave-xlate-partial-state-00 (work in progress), March 2010.

6.2. Informative References

[5] Bagnulo, M., Sullivan, A., Matthews, P., and I. Beijnum, "DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", draft-ietf-behave-dns64-00 (work in progress), July 2009.

Authors' Addresses

Charles E. Perkins
WiChorus Inc.
3590 N. 1st Street, Suite 300
San Jose CA 95134
USA

Email: charliep@computer.org

Xing Li
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
China

Email: xing@cernet.edu.cn

Congxiao Bao
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
China

Email: congxiao@cernet.edu.cn